arXiv:1410.3111v1 [stat.ML] 12 Oct 2014

Hierarchical models for neural population dynamics in the presence of non-stationarity

Mijung Park
Gatsby Computational Neuroscience Unit, University College London
Max Planck Institute for Biological Cybernetics Tübingen &
Bernstein Center for Computational Neuroscience Tübingen
mijung@gatsby.ucl.ac.uk

Jakob H. Macke
Max Planck Institute for Biological Cybernetics Tübingen
Bernstein Center for Computational Neuroscience Tübingen
jakob@tuebingen.mpg.de

Abstract

Neural population activity often exhibits rich variability and temporal structure. This variability is thought to arise from single-neuron stochasticity, neural dynamics on short time-scales, as well as from modulations of neural firing properties on long time-scales, often referred to as "non-stationarity". To better understand the nature of co-variability in neural circuits and their impact on cortical information processing, we need statistical models that are able to capture multiple sources of variability on different time-scales. Here, we introduce a hierarchical statistical model of neural population activity which models both neural population dynamics as well as inter-trial modulations in firing rates. In addition, we extend the model to allow us to capture non-stationarities in the population dynamics itself (i.e., correlations across neurons). We develop variational inference methods for learning model parameters, and demonstrate that the method can recover non-stationarities in both average firing rates and correlation structure. Applied to neural population recordings from anesthetized macaque primary visual cortex, our models provide a better account of the structure of neural firing than stationary dynamics models.

1 Introduction

Neural spiking activity recorded from populations of cortical neurons exhibits substantial variability in response to repeated presentations of a sensory stimulus [1]. This variability is thought to arise both from dynamics generated endogenously within the circuit [2], as well as from internal and behavioural states [3–5]. An understanding of how the interplay between sensory inputs and endogenous dynamics shapes neural activity patterns is essential for our understanding of how information is processed by neuronal populations. Multiple statistical [6–11] and mechanistic [12] models for characterising neuronal population dynamics have been developed.

In addition to these dynamics which take place on fast time-scales (milliseconds up to few seconds), there are also processes modulating neural firing activity which take place on much slower time-scales (seconds to days). The slow drifts in rates across an experiment can be caused by fluctuations in arousal, anaesthesia level or physiological properties of the experimental preparation [13–15]. Furthermore, processes such as learning and short-term plasticity can lead to changes in neural

firing properties [16]. The statistical structure of these slow fluctuations has been modelled using state-space models and related techniques [17–21].

To accurately capture the the structure of neural dynamics and to disentangle the contributions of slow and fast modulatory processes to neural variability and co-variability, it is therefore important to develop models that can capture neural dynamics both on fast (i.e., within experimental trials) and slow (i.e., across trials) time-scales. Czanner et al. [22] presented a statistical model of single-neuron firing in which within-trial dynamics are modelled by (generalised) linear coupling from the recent spiking history of each neuron onto its instantaneous firing rate, and across-trial dynamics were modelled by defining a random walk model over parameters. More recently, Mangion et al [23] presented a latent linear dynamical system model with Poisson observations (PLDS, [6,9,11]) with a one-dimensional latent space, and used a heuristic filtering-approach for tracking parameters, again based on a random-walk model.

Here, we present a hierarchical Bayesian model of neural dynamics. Our model consists of a latent dynamical system with Poisson observations (PLDS) to model neural population dynamics combined with a Gaussian process (GP) [24] to model modulations in firing rates or model-parameters across experimental trials. Compared to the random-walk based prior work, using a GP is a more flexible and powerful way of defining the distributions over the non-stationarity, with hyperparameters that control its variability and smoothness.

We focus on two concrete variants of this general model: In the first variant, we introduce a new set of variables which control neural firing rate on each trial to capture non-stationarity in firing rates. In the second variant, we allow the dynamics matrix which determines the spatio-temporal correlations in the population to vary across trials. We derive a variational Bayesian (VB) inference method for estimating the posterior distribution over (possibly time-varying) parameters from population recordings of spiking activity. Our approach generalises the 1-dimensional latent states in [23] to models with multi-dimensional states, as well as to a Bayesian treatment of non-stationarity based on Gaussian Process priors.

2 Hierarchical non-stationary models of neural population dynamics

We start by introducing a hierarchical model for capturing short time-scale population dynamics as well as long time-scale non-stationarities in firing rates or model parameters. Although we use the term "non-stationary" to mean that the system is best described by parameters that change over time, the distribution over parameters can be described by a stochastic process which might be strictly stationary in the statistical sense[1].

Modelling framework. We assume that the neural population activity $\mathbf{y}_t$ depends on a k-dimensional latent state $\mathbf{x}_t \in \mathbb{R}^k$ and a modulatory factor $\mathbf{h}^{(i)} \in \mathbb{R}^k$ which is different for each trial $i = \{1, \dots, r\}$. The latent state $\mathbf{x}$ models short-term co-variability of spiking activity and the modulatory factor $\mathbf{h}$ models slowly varying mean firing rates across experimental trials.

We model neural spiking activity of p neurons as conditionally Poisson given the latent state $\mathbf{x}_t$ and a modulator $\mathbf{h}^{(i)}$, with a log firing rate which is linear in parameters and latent factors,

$$\mathbf{y}_t | \mathbf{x}_t, C, \mathbf{h}^{(i)}, \mathbf{d} \sim \text{Poiss}(\mathbf{y}_t | \exp(C(\mathbf{x}_t + \mathbf{h}^{(i)}) + \mathbf{d})), \tag{1}$$

where the loading matrix $C \in \mathbb{R}^{p \times k}$ specifies how each neuron is related to the latent state and the modulator, and $\mathbf{d} \in \mathbb{R}^p$ is an offset term that controls the mean firing rate of each cell. We note that, because of the use of an exponential firing-rate nonlinearity, latent factors have a multiplicative effect on neural firing rates, as has been observed experimentally [5, 15].

Following [9, 11, 23], we assume that the latent dynamics evolve according to a first-order autoregressive process with Gaussian innovations,

$$\mathbf{x}_t | \mathbf{x}_{t-1}, A, B, Q \sim \mathcal{N}(\mathbf{x}_t | A\mathbf{x}_{t-1} + B\mathbf{u}_t, Q). \tag{2}$$

Here, we allow for sensory stimuli (or experimental covariates), $\mathbf{u}_t \in \mathbb{R}^d$ to influence the latent states linearly. The dynamics matrix $A \in \mathbb{R}^{k \times k}$ determines the state evolution, $B \in \mathbb{R}^{k \times d}$ models

[1] A stochastic process is strict-sense stationary if its joint distribution over any two time-points t and s only depends on the elapsed time $t - s$.

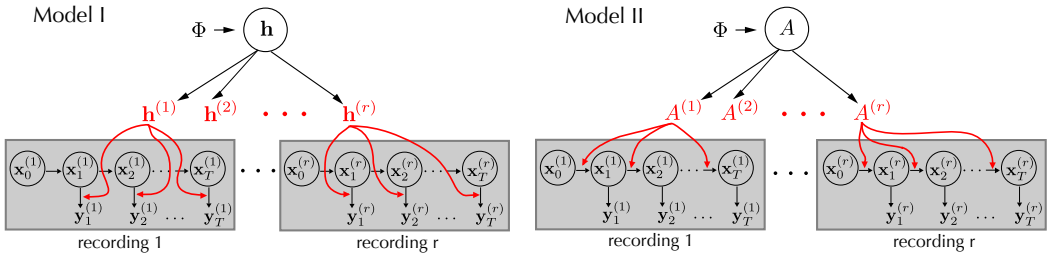

Figure 1. Schematic of hierarchical non-stationary Poisson observation Latent Dynamical System (N-PLDS) models. **Model I** for capturing non-stationarity in mean firing rates. The parameter $\mathbf{h}$ slowly varies across trials and leads to fluctuations in mean firing rates. **Model II** for capturing non-stationarity in population dynamics. The dynamics matrix A changes across trials, as controlled by the hyperparameters Φ.

the dependence of latent states on external inputs, and $Q \in \mathbb{R}^{k\times k}$ is the noise covariance $Q \in \mathbb{R}^{k\times k}$. We set Q to be the identity matrix, $Q = \mathbf{I}_k$ as in [25], and we assume $\mathbf{x}_0^{(i)} \sim \mathcal{N}(0, \mathbf{I}_k)$.

The parameters in this model are $\boldsymbol{\theta} = \{A, B, C, \mathbf{d}, \mathbf{h}^{(1:r)}\}$. We refer to this general model as *non-stationary PLDS* (N-PLDS). Different variants of N-PLDS can be constructed by placing priors on individual parameters which allow them to vary across trials (in which case they would then depend on the trial index i) or by omitting different components of the model. In the following, we will focus on two concrete variants.

Model I : Non-stationarity in mean firing rates. In the first model (Fig. 1 Left), we assume all parameters apart from $\mathbf{h}$ to be constant across trials. For the modulator $\mathbf{h}$, we assume that it varies across trials according to a GP with mean 0 and (modified) squared exponential kernel, i.e. $\mathbf{h}^{(i)} \sim \mathcal{GP}(0, K(i,j))$, where the (i,j)th block of K (size $k \times k$) is given by

$$K(i,j) = (\sigma^2 + \epsilon\delta_{i,j}) \exp\left(-\tfrac{1}{2\tau^2}(i-j)^2\right) I_k. \tag{3}$$

Here, we assume the independent noise-variance on the diagonal (ϵ) to be constant and small as in [26]. When $\sigma^2 = \epsilon = 0$, the modulator vanishes, which corresponds to the conventional PLDS model with fixed parameters [9, 11]. When $\sigma^2 > 0$, the mean firing rates vary across trials, and the parameter τ determines the time-scale (in units of 'trials') of these fluctuations. We impose ridge priors on the model parameters (see Appendix A for details), so that the total set of hyperparameters of the model is $\Phi = \{\sigma^2, \tau^2, \boldsymbol{\phi}\}$, where $\boldsymbol{\phi}$ is the set of ridge parameters.

Model II : Non-stationarity in population dynamics. With the second variant (Fig. 1 Right), we aim to capture non-stationarity in dynamics. We drop the modulator $\mathbf{h}$, and instead allow the dynamics matrix A to vary across trials. That is, we allow the model to have a different matrix $A^{(i)}$ for each trial i. We assume that the evolution of A's over time follows a Gaussian process, $\mathbf{a}^{(i)} \sim \mathcal{GP}(\bar{\mathbf{a}}, K(i,j))$, defined over the vectorized matrices $\mathbf{a}^{(i)} = \text{vec}(A^{(i)\top}) \in \mathbb{R}^{k^2}$. Here, $\bar{\mathbf{a}}$ is the (vectorized) mean dynamics matrix, the covariance kernel K is as for model I, with the only exception that the blocks are now of size $k^2 \times k^2$ as each $\mathbf{a}^{(i)}$ has k^2 elements. As before, when $\sigma^2 = \epsilon = 0$, this model reduces to a stationary PLDS. When τ goes to zero, this model corresponds to independently drawing $\mathbf{a}^{(i)}$ from $\mathcal{N}(\bar{\mathbf{a}}, \sigma^2 I_{k^2})$, i.e. a model in which the dynamics are slightly different on each trial, but the fluctuations are not correlated across trials. For larger values of τ, the squared exponential kernel ensures that A will vary smoothly across trials. The parameters in this model are $\boldsymbol{\theta} = \{A^{(1:r)}, B, C, \mathbf{d}\}$. We also impose ridge priors on $\{B, C, \mathbf{d}\}$, which results in hyperparameters $\Phi = \{\sigma^2, \tau^2, \boldsymbol{\phi}\}$.

3 Variational Bayesian Expectation Maximization

Our goal is to infer parameters and latent variables in each model. The exact posterior distribution in each model is not analytically tractable due to the Poisson likelihood term. Here, we approximate the posterior over the parameters and the latent variables by maximising the lower bound of the

marginal likelihood of the observations,

$$\log p(\mathbf{y}_{1:T}^{(1:r)}) \geq \int d\boldsymbol{\theta}\; d\mathbf{x}_{1:T}^{(1:r)}\; q(\boldsymbol{\theta}, \mathbf{x}_{1:T}^{(1:r)})\; \log \frac{p(\boldsymbol{\theta}, \mathbf{x}_{1:T}^{(1:r)}, \mathbf{y}_{1:T}^{(1:r)})}{q(\boldsymbol{\theta}, \mathbf{x}_{1:T}^{(1:r)})}, \tag{4}$$

where the approximate posterior is denoted by $q(\boldsymbol{\theta}, \mathbf{x}_{1:T}^{(1:r)}) = q_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \prod_{i=1}^{r} q_{\mathbf{x}}(\mathbf{x}_{0:T}^{(i)})$, which we assume to be factorized between latents and parameters. We maximize the lower bound by iterating the variational Bayesian expectation maximization (VBEM) algorithm [25], which consists of (1) the variational Bayesian expectation (VBE) step for computing $q_{\mathbf{x}}(\mathbf{x}_{0:T}^{(1:r)})$,

$$q_{\mathbf{x}}(\mathbf{x}_{0:T}^{(1:r)}) \propto \exp\left[\int d\boldsymbol{\theta} q_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \log p(\mathbf{x}_{1:T}^{(1:r)}, \mathbf{y}_{1:T}^{(1:r)} | \boldsymbol{\theta})\right], \tag{5}$$

and (2) the variational Bayesian maximization (VBM) step for computing $q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$,

$$q_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \propto p(\boldsymbol{\theta}) \exp\left[\int d\mathbf{x}_{0:T}^{(1:r)} q_{\mathbf{x}}(\mathbf{x}_{0:T}^{(1:r)}) \log p(\mathbf{x}_{0:T}^{(1:r)}, \mathbf{y}_{1:T}^{(1:r)} | \boldsymbol{\theta})\right]. \tag{6}$$

While a complete derivation of the algorithm is presented in Appendix A, we here highlight its main steps.

VBE step. Using the first-order dependency in latent states, we derive a sequential forward/backward algorithm to obtain $q_{\mathbf{x}}(\mathbf{x}_{0:T}^{(1:r)})$, generalising the approach of [23] to multi-dimensional latent states. Since the VBE step decouples across trials, this step is easy to parallelise across trials (and is not more expensive than Bayesian inference for a 'fixed parameter' PLDS) and we omit the trial-indices for clarity. The forward message $\alpha(\mathbf{x}_t)$ at time t is given by

$$\alpha(\mathbf{x}_t) \propto \int d\mathbf{x}_{t-1} \alpha(\mathbf{x}_{t-1}) \exp\left[\langle \log(p(\mathbf{x}_t|\mathbf{x}_{t-1}) p(\mathbf{y}_t|\mathbf{x}_t)) \rangle_{q_{\boldsymbol{\theta}}(\boldsymbol{\theta})}\right]. \tag{7}$$

Assuming that the forward message at time $t-1$ denoted by $\alpha(\mathbf{x}_{t-1})$ is Gaussian, the Poisson likelihood term will render the forward message at time t non-Gaussian, but we will approximate $\alpha(\mathbf{x}_t)$ as a Gaussian using the first and second derivatives of the right-hand side of eq. 7. [2]

Similarly, the backward message at time $t-1$ is given by

$$\beta(\mathbf{x}_{t-1}) \propto \int d\mathbf{x}_t \beta(\mathbf{x}_t) \exp\left(\langle \log(p(\mathbf{x}_t|\mathbf{x}_{t-1}) p(\mathbf{y}_t|\mathbf{x}_t)) \rangle_{q_{\boldsymbol{\theta}}(\boldsymbol{\theta})}\right), \tag{8}$$

which we also approximate to a Gaussian for tractability in computing backward messages.

Using the forward/backward messages, we compute the posterior marginal distribution over latent variables (See Appendix A). We also need to compute the cross-covariance between neighbouring latent variables to obtain the sufficient statistics of latent variables (which we will need for VBM step). The pairwise marginals of latent variables are given by

$$p(\mathbf{x}_t, \mathbf{x}_{t+1}|\mathbf{y}_{1:T}) \quad \propto \quad \beta(\mathbf{x}_{t+1}) \exp\left(\langle \log(p(\mathbf{y}_{t+1}|\mathbf{x}_{t+1}) p(\mathbf{x}_{t+1}|\mathbf{x}_t)) \rangle_{q_{\boldsymbol{\theta}}(\boldsymbol{\theta})}\right) \alpha(\mathbf{x}_t), \tag{9}$$

which we approximate as a jointly Gaussian distribution by using the first/second derivatives of eq. 9 and extracting the cross-covariance term from the joint covariance matrix.

VBM step. In the VBM step, we need to update the posterior distribution over the parameters. Under Model I, the posterior over parameters factorizes as

$$q_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \quad = \quad q_{\mathbf{a},\mathbf{b}}(\mathbf{a}, \mathbf{b})\; q_{\mathbf{c},\mathbf{d},\mathbf{h}}(\mathbf{c}, \mathbf{d}, \mathbf{h}^{(1:r)}), \tag{10}$$

where vectorized notations $\mathbf{b} = \mathrm{vec}(B^\top)$ and $\mathbf{c} = \mathrm{vec}(C^\top)$. For this model, we set $\mathbf{c}, \mathbf{d}$ to the maximum likelihood estimates $\hat{\mathbf{c}}, \hat{\mathbf{d}}$ for simplicity in inference. The computational cost of this algorithm is dominated by the cost of calculating the posterior distribution over $\mathbf{h}^{(1:r)}$, which involves manipulation of a rk-dimensional Gaussian. While this was still tractable without further approximations for the data-set sizes used in our analyses below, we note that a variety of approximate methods for GP-inference exist which could be used to improve efficiency of this computation. In particular, we will typically be dealing with systems in which $\tau \gg 1$, which means that the kernel-matrix is smooth and can be approximated using low-rank representations [24].

Under Model II, the posterior over $\mathbf{a}$ is replaced by $\mathbf{a}^{(1:r)}$, and $\mathbf{h}^{(1:r)}$ is omitted from eq. 10. We then approximate each factor in eq. 10 to a multivariate normal distribution using the first/second derivatives of eq. 6 w.r.t. each of the parameters. Again, computation time is dominated be the GP term, which however lends itself well to low-rank approximations.

[2] This approximation does not guarantee monotonically increasing lower bound after each VBEM iteration. However, we monitored its convergence by computing one-step ahead prediction scores and found that the approximation yielded accurate results as shown in Applications.

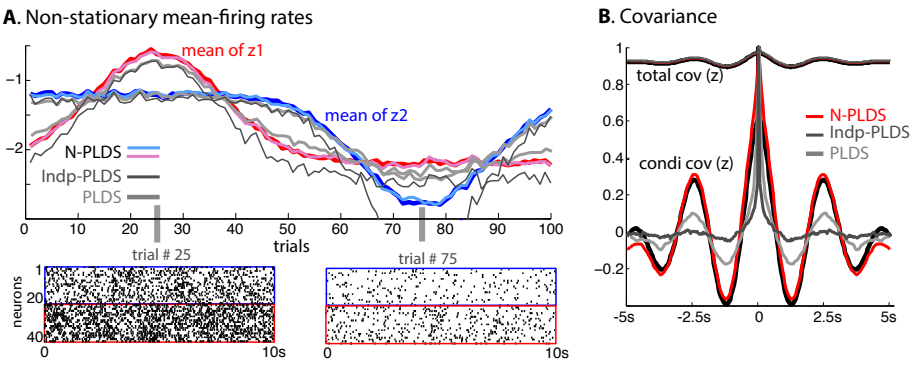

Figure 2. Illustration of non-stationarity in firing rates (data simulated from Model I). **A**: Spike rates of 40 neurons are influenced by two slowly varying firing rate modulators. The log mean firing rates of the two groups of neurons are z_1(red, group 1) and z_2(blue, group 2) across 100 trials (top). Raster plots (bottom) show the extreme cases, i.e. trials 25 and 75. The traces show the posterior mean of z estimated by N-PLDS (Model I, light blue for z_2, light red for z_1), independent PLDSs (fit a PLDS to each trial data individually, dark gray), and PLDS (light gray). **B**: Total and conditional (on each trial) covariance of recovered neural responses from each model (averaged across all neuron pairs, and then normalised for visualisation). The covariances recovered by our model (red) well match the true ones (black), while those by independent PLDSs (gray) and a single PLDS (light gray) do not.

Hyperparameter estimation. In each iteration of VBEM algorithm, we also update the hyperparameters by maximizing the lower bound w.r.t. hyperparameters. Given the approximate posterior $q(\boldsymbol{\theta}, \mathbf{x}_{0:T}^{(1:r)})$, the variational lower bound simplifies to (see Ch.5 in [25] for details)

$$\log p(\mathbf{y}_{1:T}^{(1:r)}|\Phi) \geq -KL(\Phi) + c. \tag{11}$$

The KL divergence between the prior and posterior over parameters, denoted by $\mathcal{N}(\boldsymbol{\mu}_\Phi, \Sigma_\Phi)$ and $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$, respectively, is given by

$$KL(\Phi) = -\tfrac{1}{2}\log|\Sigma_\Phi^{-1}\Sigma| + \tfrac{1}{2}\mathrm{Tr}\left[\Sigma_\Phi^{-1}\Sigma\right] + \tfrac{1}{2}(\boldsymbol{\mu}-\boldsymbol{\mu}_\Phi)^\top\Sigma_\Phi^{-1}(\boldsymbol{\mu}-\boldsymbol{\mu}_\Phi) + c, \tag{12}$$

where the prior mean and covariance depend on the hyperparameters.

Predictive distributions for test data. In our model, different trials are no longer considered as independent, so we can predict parameters for held-out trials. Using the GP model on $\mathbf{h}$ and $\mathbf{a}$, we have closed form predictive distributions on $\mathbf{h}^*$ and $\mathbf{a}^*$ for test data $\mathcal{D}^*$ given training data $\mathcal{D}$:

$$p(\mathbf{h}^*|\mathcal{D},\mathcal{D}^*) = \mathcal{N}(K^*K^{-1}\boldsymbol{\mu}_\mathbf{h},\ K^{**} - K^*(K+H_\mathbf{h}^{-1})^{-1}K^{*\top}), \tag{13}$$

$$p(\mathbf{a}^*|\mathcal{D},\mathcal{D}^*) = \mathcal{N}(\bar{\mathbf{a}} + K^*K^{-1}(\boldsymbol{\mu}_\mathbf{a}-\bar{\mathbf{a}}), K^{**} - K^*(K+H_\mathbf{a}^{-1})^{-1}K^{*\top}), \tag{14}$$

where K is the prior covariance matrix on $\mathcal{D}$ and K^{**} is on $\mathcal{D}^*$, and K^* is their prior cross-covariance as introduced in Ch.2 of [24], and the negative Hessians $H_\mathbf{h}$ and $H_\mathbf{a}$ are defined in Appendix A.

4 Applications

4.1 Simulated data

We first illustrate the performance of Model I on a simulated (from Model I) population recording from 40 neurons consisting of 100 trials of length $T = 200$ time steps each. We used a 4-dimensional latent state and assumed that the population consisted of two homogeneous sub-populations of size 20 each, with one modulatory input controlling rate fluctuations in each group (See Fig. 2 A). In addition, we assumed that for half of each trial, there was a time-varying stimulus ('drifting grating'), represented by a 3-dimensional vector which consisted of the sine and cosine of the time-varying phase of the stimulus (frequency 0.4 Hz) as well as an additional term which indicated whether the stimulus was on.

We fit Model-I N-PLDS to the data, and found that it successfully captures the non-stationarity in (log) mean firing rates, defined by $z = C(\mathbf{x}+\mathbf{h}) + \mathbf{d}$, as shown in Fig. 2 A, and recovers the total

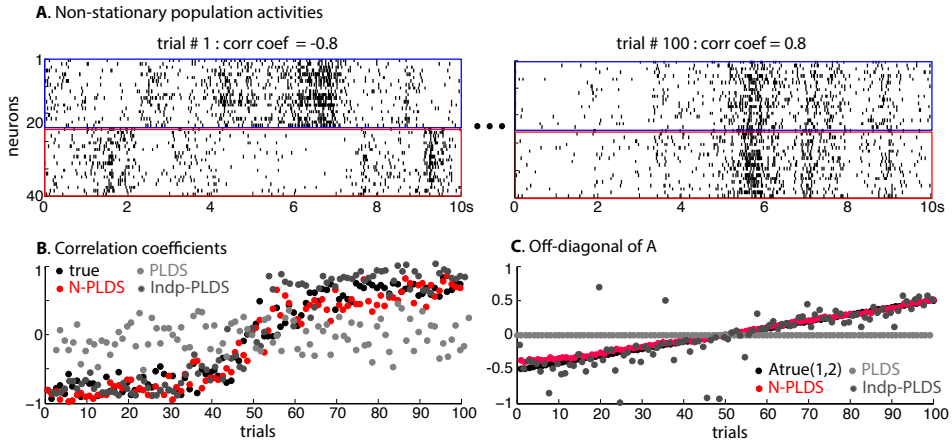

Figure 3. Illustration of non-stationarity in population dynamics (data simulated from Model II). A: Raster plots of spontaneous activity from 40 neurons during 10 seconds of recording for simulated trials 1 and 100. We assumed that the two sub-populations (blue and red) have negative correlation at trial 1 and positive correlation at trial 100. **B**: Recovered correlations. Our Model II (red) accurately recovers the correlations between two groups across trials (RMSE: 0.04), while other methods perform poorly: independent PLDSs fit to each trial individually give noisy results (RMSE 0.06) and a single PLDS fit across all trials cannot capture the change in correlation (RMSE 0.44). **C**: Estimation of off-diagonal in dynamics matrices. We fixed the loading matrix C to its true value to avoid issues with non-identifiability of parameters in LDS models. The off-diagonal term A_{12} estimated by our model matched the true values well, whereas the independent PLDS produced noisy estimates, and the fixed PLDS cannot capture the change in A_{12}.

and trial-conditioned covariances (the across-trial mean of the single-trial covariances of z). For comparison, we also fit 100 separate PLDSs to the data from each trial, as well as a single PLDS to the entire data. The naive approach of fitting an individual PLDS to each trial can, in principle, follow the modulation. However, as each model is only fit to one trial, the parameter-estimates are very noisy since they are not well constrained by the data from each trial. We note that a single PLDS with fixed parameters is able to track the modulations in firing rates in the posterior mean here– however, a single PLDS would not be able to extrapolate firing rates for unseen trials (as we will demonstrate in our analyses on neural data below). In addition, it will also fail to separate 'slow' and 'fast' modulations into different parameters. By comparing the total covariance of the data (averaged across neuron pairs) to the 'trial-conditioned' covariance (calculated by estimating the covariance on each trial individually, and averaging covariances) one can calculate how much of the cross-neuron co-variability can be explained by across-trials fluctuations in firing rates (see e.g., [15]). In this simulation (which illustrates an extreme case dominated by strong across-trial effects), the conditional covariance is much smaller than the true one.

Next, we tested Model II using a simulation of spontaneous activity from a population of 40 neurons (simulated from Model II). We again assumed that the population could be split into two sub-populations of size 20 neurons each, and simulated an experiment in which the correlation across the two sub-populations changed dramatically across the experiment: Specifically, we generated a 2-d latent state that controls correlations in firing rates between the two groups of neurons, and adjusted the off-diagonal term in the dynamics matrix (A_{12}) such that the correlation between the groups varied slowly from -1 to 1 across 100 trials, where the length of each trial is $T = 200$. Other elements of A were adjusted such that the stationary covariance of the system was kept constant.

As before, we fit Model II N-PLDS, a single PLDS, and 100 independent PLDSs to the data. Our model accurately recovered the correlation change in z across trials, while the single PLDS was not able to capture the non-stationarity and the independent PLDSs exhibited noisy correlations (Fig. 3). Finally, our model also accurately recovered the off-diagonal parameter A_{12} (Fig. 3 C). For panel C only, we set the loading matrix C to the ground truth value for each of the models (Model II, fixed PLDS, separate PLDSs). LDS models suffer from non-identifiability of parameters, implying that estimated parameters do not necessarily match the true parameters even for perfect model fits.

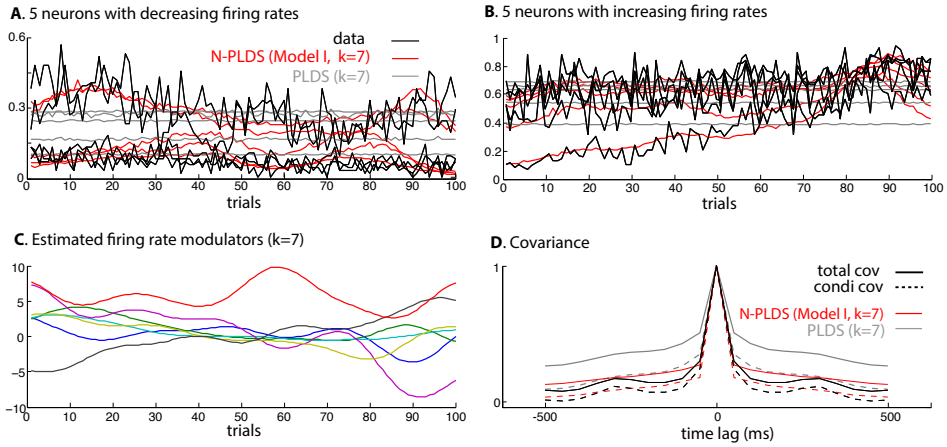

Figure 4. Non-stationary firing rates in a population of V1 neurons. A: Firing rate change of 5 neurons (black) among 64 with the strongest net-reduction in firing rate (quantified using linear regression on trial indices). **B**: Firing rate change of 5 neurons with strongest increase in firing rates. The best RMSE on predicting mean firing rates by our method was when $k = 7$, $RMSE = 0.7383$. For comparison, single fixed-parameter PLDS across trials has an RMSE of 1.2496. **C**: Estimated 7-dimensional modulator h. The modulator with an estimated length scale of ≈ 10 is smoothly varying across trials. **D**: Comparison of total and (trial) conditional covariances. Our method recovers both covariances more accurately than PLDS.

4.2 Neurophysiological data

Dataset I: Non-stationary mean firing rates

How big are non-stationarities in neural population recordings, and can our model successfully capture them? To address these questions, we first analysed a population recording from anesthetized macaque primary visual cortex consisting of 64 neurons stimulated by sine grating stimuli. The details of data collection are described in [5], but our dataset also included units that were not used in the original study due to strong non-stationarities in their firing rate. We binned the spikes recorded during 100 trials of the same orientation using 50ms bins, resulting in trials of length $T = 80$ bins. Analogously to the simulated dataset above, we parameterised the stimulus as a 3-d vector of the sine and cosine values of the phase of the grating as well as an indicator that specifies whether there is a stimulus or not. The stimulus was presented for 2 seconds. Before and after the stimulus presentation, a blank screen was presented for 1 second. We divided the data into test data (every 10th trial) and training data (i.e., the remaining 90 trials).

We then fit our N-PLDS Model I to the training data. Using the estimated parameters from the training data, we made predictions on the modulator $\mathbf{h}$ on test data (we used the mean of the predictive distribution, see methods). Using these parameters, we drew samples for spikes for the entire trials to compute the mean firing rates of each neuron at each trial. For comparison, we also fit a single PLDS to the data. As this model does not allow for across-trial modulations of firing rates, we simple kept the parameters estimated from the training data.

For visualisation of results, we used linear regression on each neuron (regression from trial index to mean firing rates) to classify neurons as 'decreasing in firing rate', 'increasing in firing rate' or 'not being modulated by firing rate' based on how much of their total variance was explained by the regression, and based on the sign of the regression coefficient. For the neurons with strongest modulations in firing rates, the modulations were well captured by our model (Fig. 4 A/B, results are for latent dimensionality $k = 7$). For non-modulated neurons (36 out of 64), our model likewise identified firing rates to be constant. We tested different latent dimensions ($k = 1, 5, 7, 9$), and found that our model works the best when $k = 7$ in terms of the RMSE on mean prediction: $1.1149, 0.9398, 0.7383, 0.7893$ for $k = 1, 5, 7, 9$, respectively. Our model successfully captures the non-stationarity in mean firing rates (Fig. 4 A/B) and also accurately recovers the total and conditional covariances (averaged across all neurons). In contrast, the fixed-parameter PLDS has firing rates which are constant across trials, therefore achieves worse reconstruction RMSE for each k (e.g. RMSE: 1.2496 for $k = 7$), and also overestimates the conditional covariance.

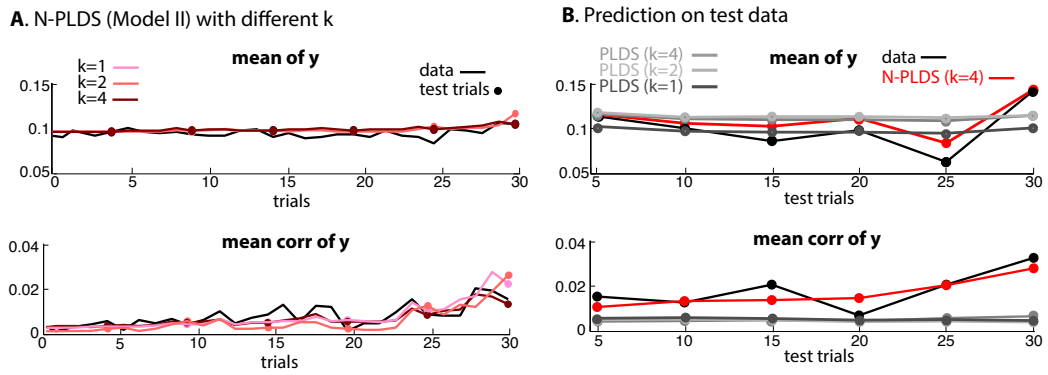

Figure 5. Non-stationary population dynamics (data from [27]). **A**: Summary statistics of samples from N-PLDS (Model II) with non-stationarity dynamics matrix A for different dimensions of latent dynamics ($k = 1, 2, 4$). The top plot shows the mean firing rate of 40 neurons during 30 epochs, showing that there is only a slight systematic drift in mean firing rate. Each dot represents predicted mean firing rates for the held-out data (6 trials). The bottom plot shows the mean correlation of the spike counts. All three N-PLDS models capture the increase in correlation at the end of the experiment, with the $k = 4$ capturing it most accurately. **B**: Comparison to using a PLDS model with fixed parameters ($k = 1, 2, 4$). Both the mean firing rate and correlation in PLDS are constant across epochs. As a consequence, the best RMSE on mean correlation estimation in PLDS is 0.0138 ($k = 1$) compared to 0.0087 ($k = 4$) in N-PLDS.

Dataset II : Non-stationary population dynamics

Finally, we analyzed a dataset of spontaneous activity recorded from a population of 40 neurons from macaque visual cortex. The details of data collection are described in [27] and the data is available from [28]. Using the spike-sorting information provided in the dataset, we selected the spike-cluster with highest signal-to-noise ratio from each recording channel, and out of those 46 units kept the 40 units with highest firing rates. As the original data consisted of one continuous recording of length 15 minutes, we divided the data into 30 'epochs' of length 30 seconds each, and used every 5th epoch (20% of the data) for testing and the rest (80% of data) for training.

In this data, the mean firing rates are almost constant across time, while the correlations increase at the end of the experiment (Fig. 5 A). After estimating the parameters of our N-PLDS (Model II) from the training data, we computed the predictive distribution on the dynamics matrices A^* for the test data. Using these parameters, we drew samples for spikes to compute the mean firing rates for each trial (Fig. 5 A), as well as the mean pairwise cross-correlations across all neuron pairs. The correlations estimated from N-PLDS (Model II) matched those in the data. For PLDS with fixed parameters, the estimated firing rates and correlations are constant across epochs (Fig. 5 B). To quantify these results, we computed the RMSE in the prediction of mean firing rates and mean correlations on test epochs. The RMSEs on mean firing rate estimation for PLDS are $0.0156, 0.0182, 0.0188$ for $k = 1, 2, 4$, respectively, while RMSE of N-PLDS is 0.0080 ($k = 4$). The RMSE on mean correlation estimation in PLDSs is 0.0138 (same for $k = 1, 2, 4$) and 0.0087 ($k = 4$) in N-PLDS.

5 Discussion

Non-stationarities are ubiquitous in neural data: Slow modulations in firing properties can result from diverse processes such as plasticity and learning, fluctuations in arousal, cortical reorganisation after injury as well as development and aging. In addition, non-stationarities in neural data can also be a consequence of experimental artifacts, and can be caused by fluctuations in anaesthesia level, stability of the physiological preparation or electrode drift. Whatever the the origins of non-stationarities are, it is important to have statistical models which can identify them and disentangle their effects from correlations and dynamics on faster time-scales [14].

We here presented a hierarchical Bayesian model for neural population dynamics in the presence of non-stationarity. Specifically, we concentrated on two variants of this model which focus on non-stationarity in firing rates (Model I) and dynamics matrices (Model II). Applied to two sets

of neurophysiological recording data, we demonstrated that this modelling approach can successfully capture these types of non-stationarities in neurophysiological recordings from primary visual cortex.

There are limitations to the current study: (1) how to select amongst multiple different models which could be used to model neural non-stationarity for a given dataset; (2) how to scale up the current algorithm for larger trial numbers (e.g., using low-rank approximations to the covariance matrix); and (3) how to overcome the slow convergence properties of GP kernel parameter estimation [29]. We believe that extending our method to address these questions provides an exciting direction for future research, and will result in a powerful set of statistical methods for investigating how neural systems operate in the presence of non-stationarity.

Acknowledgments

We thank the authors of [5] for sharing dataset I with us and making it publicly available at: http://toliaslab.org/publications/ecker-et-al-2014/, as well as the authors of [27] for making dataset II publicly available via CRCNS.org. We thank Srinivas Turaga, Thomas Desautels, Alexander Ecker, and Maneesh Sahani for their valuable comments on the manuscript. This work was funded by the German Federal Ministry of Education and Research (BMBF; FKZ: 01GQ1002, Bernstein Center Tübingen) and the Gatsby Charitable Trust.

References

[1] A. Renart and C. K. Machens. Variability in neural activity and behavior. *Curr Opin Neurobiol*, 25:211–20, Apr 2014.

[2] A. Destexhe. Intracellular and computational evidence for a dominant role of internal network activity in cortical computations. *Curr Opin Neurobiol*, 21(5):717–725, 2011.

[3] G. Maimon. Modulation of visual physiology by behavioral state in monkeys, mice, and flies. *Curr Opin Neurobiol*, 21(4):559–64, 2011.

[4] K. D. Harris and A. Thiele. Cortical state and attention. *Nat Rev Neurosci*, 12(9):509–523, 2011.

[5] A. S. Ecker, P. Berens, R. J. Cotton, M. Subramaniyan, G. H. Denfield, C. R. Cadwell, S. M. Smirnakis, M. Bethge, and A. S. Tolias. State dependence of noise correlations in macaque primary visual cortex. *Neuron*, 82(1):235–48, 2014.

[6] A. C. Smith and E. N. Brown. Estimating a state-space model from point process observations. *Neural Comput*, 15(5):965–91, 2003.

[7] U. T. Eden, L. M. Frank, R. Barbieri, V. Solo, and E. N. Brown. Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Comput*, 16(5):971–98, 2004.

[8] B. M. Yu, A. Afshar, G. Santhanam, S. I. Ryu, K. Shenoy, and M. Sahani. Extracting dynamical structure embedded in neural activity. In *NIPS 18*, pages 1545–1552. MIT Press, Cambridge, MA, 2006.

[9] J. E. Kulkarni and L. Paninski. Common-input models for multiple neural spike-train data. *Network*, 18(4):375–407, 2007.

[10] W. Truccolo, L. R. Hochberg, and J. P. Donoghue. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat Neurosci*, 13(1):105–111, 2010.

[11] J. H. Macke, L. Buesing, J. P. Cunningham, B. M. Yu, K. V. Shenoy, and M. Sahani. Empirical models of spiking in neural populations. In *NIPS*, pages 1350–1358, 2011.

[12] C. van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–6, Dec 1996.

[13] G. J. Tomko and D. R. Crapper. Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Res*, 79(3):405–18, Oct 1974.

[14] C. D. Brody. Correlations without synchrony. *Neural Comput*, 11(7):1537–51, 1999.

[15] R. L. T. Goris, J. A. Movshon, and E. P. Simoncelli. Partitioning neuronal variability. *Nat Neurosci*, 17(6):858–65, Jun 2014.

[16] C. D. Gilbert and W. Li. Adult visual cortical plasticity. *Neuron*, 75(2):250–64, Jul 2012.

[17] E. N. Brown, D. P. Nguyen, L. M. Frank, M. A. Wilson, and V. Solo. An analysis of neural receptive field plasticity by point process adaptive filtering. *Proc Natl Acad Sci U S A*, 98(21):12261–6, 2001.

[18] L. M. Frank, U. T. Eden, V. Solo, M. A. Wilson, and E. N. Brown. Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: an adaptive filtering approach. *J Neurosci*, 22(9):3817–30, May 2002.

[19] N. A. Lesica and G. B. Stanley. Improved tracking of time-varying encoding properties of visual neurons by extended recursive least-squares. *IEEE Trans Neural Syst Rehabil Eng*, 13(2):194–200, Jun 2005.

[20] V. Ventura, C. Cai, and R.E. Kass. Trial-to-Trial Variability and Its Effect on Time-Varying Dependency Between Two Neurons. *Journal of Neurophysiology*, 94(4):2928–2939, 2005.

[21] C. S. Quiroga-Lombard, J. Hass, and D. Durstewitz. Method for stationarity-segmentation of spike train data with application to the pearson cross-correlation. *J Neurophysiol*, 110(2):562–72, Jul 2013.

[22] Gabriela C., Uri T. E., Sylvia W., Marianna Y., Wendy A. S., and Emery N. B. Analysis of between-trial and within-trial neural spiking dynamics. *Journal of Neurophysiology*, 99(5):2672–2693, 2008.

[23] A. Z. Mangion, K. Yuan, V. Kadirkamanathan, M. Niranjan, and G. Sanguinetti. Online variational inference for state-space models with point-process observations. *Neural Comput*, 23(8):1967–1999, 2011.

[24] C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*. MIT Press Cambridge, MA, USA, 2006.

[25] M. J. Beal. *Variational Algorithms for Approximate Bayesian Inference*. PhD thesis, Gatsby Unit, University College London, 2003.

[26] B. M. Yu, J. P. Cunningham, G. Santhanam, S. I. Ryu, K. V. Shenoy, and M. Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, 102(1):614–635, 2009.

[27] Cheng C.J. Chu, Ping F. Chien, and Chou P. Hung. Tuning dissimilarity explains short distance decline of spontaneous spike correlation in macaque {V1}. *Vision Research*, 96(0):113 – 132, 2014.

[28] Chou P. Hung Cheng C. J. Chu, Ping F. Chien. Multi-electrode recordings of ongoing activity and responses to parametric stimuli in macaque v1. CRCNS.org., 2014.

[29] I. Murray and R. P. Adams. Slice sampling covariance hyperparameters of latent Gaussian models. In *NIPS 23*, pages 1723–1731. 2010.